FEMOSAA: Feature Guided and Knee Driven Multi-Objective Optimization for Self-Adaptive Software at Runtime

Tao Chen, Ke Li, Rami Bahsoon, Xin Yao
Centre of Excellence for Research in Computational Intelligence and Applications
School of Computer Science, University of Birmingham
Edgbaston, Birmingham, UK, B15 2TT
{t.chen, k.li.1, r.bahsoon, x.yao}@cs.bham.ac.uk

arXiv:1608.08933v1 [cs.SE] 31 Aug 2016

***Abstract*—Self-adaptive software (SAS) can reconfigure itself to adapt to the changing environment at runtime. Such a behavior permits continual optimization on various conflicting non-functional objectives, e.g., response time and energy consumption. In this paper, we present FEMOSAA, a novel framework that automatically synergizes the feature model and Multi-Objective Evolutionary Algorithm (MOEA), to optimize SAS at runtime. At design time, FEMOSAA automatically transposes the design of SAS, which is expressed as a feature model, to the chromosome representation and the reproduction operators (mutation and crossover) in MOEA. At runtime, the feature model serves as the domain knowledge to guide the search, providing more chances to find better solutions. In addition, we have designed a new method to search for the knee solutions, which can achieve balanced trade-off. We experimentally compare FEMOSAA with different variants and state-of-the-art approaches on a real world SAS. The results reveal its effectiveness and superiority over the others.**

***Keywords*-search-based software engineering; feature model; multi-objective evolutionary algorithm; self-adaptive software;**

I. Introduction

Self-adaptive software (SAS) is a special type of software that is capable to adapt and reconfigure itself at runtime, through a set of known features (e.g., CPU cap, thread pool size and cache size), according to the changing environment [1]. One of the goals for SAS is to continually optimize multiple and often conflicting non-functional objectives, e.g., *response time* and *energy consumption*. However, given the dynamic and uncertain nature of running software, it is difficult to fully specify all possible conditions and their adaptation solutions at design time. Thus, designing an efficient and effective runtime optimization approach is necessary, yet challenging. Depending on the complexity of SAS, it might be possible for the software engineers to design methods that enumerate all combinations of configurations at runtime. However, modern SAS often exhibits high variability, leading to explosion of the search space of all possible solutions; henceforth, enumeration can fail to scale at runtime. As a result, search-based software engineering techniques, particularly Evolutionary Algorithms (EAs), are appealing for SAS optimization [2].

To optimize SAS at runtime using EA, there are two major challenges: (i) it is difficult to convert the SAS design to the EA's context while considering the necessary features and their dependencies, e.g., one can change *Cache Mode* only if the *Cache* feature is 'turned on'. Dependency becomes even more complex in the presence of numeric features, e.g. in Tomcat [3], the size of *maxThreads* should not be less than the size of *minSpareThreads*; (ii) optimizing multiple conflicting objectives and managing their trade-offs are complex and challenging.

Most existing work fails to handle the first challenge as they have relied on a manual and/or incomplete conversion of the SAS design into the EA's context [4]–[7], which renders the process expensive, non-systematic and error-prone. Moreover, the feature dependencies are often ignored, wasting valuable function evaluations on invalid solutions while providing no guarantee on finding the valid ones. Inspired by the applications of EA to Software Product Line problems [8], researchers [9] [10] have combined the feature model [11] with the search-based techniques for optimizing SAS at runtime. However, they did not explicitly consider the numeric features and they encoded all the features into a simple binary representation to represent a solution. This might lead to the *curse of dimensionality* in the optimization, entailing unnecessary complexity. Further, they cannot prevent wasteful exploration of invalid solutions and consider very limited types of dependency.

For the second challenge above, many existing approaches simply transform the multi-objective problem into a single-objective one via aggregation, e.g., a weighted sum, before using an EA (e.g., Genetic Algorithm) [4] [12] or exact search [9] [13]. However, assigning weights to different objectives is a non-trivial task for software engineers and such an aggregation can hardly maintain the solution diversity. To alleviate this issue, studies [5]–[7] have used NSGA-II [14], a popular Multi-Objective EA (MOEA), to optimize SAS without any weights. However, NSGA-II has a coarse diversity preservation mechanism that is unable to provide well distributed solutions [15]. In addition, given the fact that MOEAs produce a set of non-dominated solutions, there is no established method for the SAS to choose an appropriate one for adaptation, entailing the risk of imbalanced trade-offs.

To address the above challenges and limitations, this paper presents FEMOSAA, a novel framework that automatically

synergizes feature model and EA/MOEA, to optimize SAS at runtime. Specifically, our novel contributions include:

– We use the feature model to represent SAS with explicit considerations of numeric features and their dependencies. In FEMOSAA, we provide an automatic approach to transpose this model into EA/MOEA's context at design time.
– To reduce the search space, our approach identifies the elitist features from the feature model. These features form the chromosome representation in EA/MOEA without destroying the original variability.
– To guide the search and avoid exploring invalid solutions, our approach extracts the feature dependencies w.r.t. the elitist features. Then, dependencies are injected into the mutation and crossover operators of the EA/MOEA.
– Without loss of generality, we exploit MOEA based Decomposition with STable-Matching model (MOEA/D-STM) [16], an algorithm that produces better convergence and diversity than NSGA-II [15], as the MOEA in FEMOSAA for optimizing SAS at runtime.
– To achieve balanced trade-off in SAS optimization, FEMOSAA finds knee solutions from the final non-dominated set. These solutions often imply well balanced trade-offs, such that any improvement on one objective of a knee will cause severe degradations on others.
– We conduct comprehensive experiments on a SAS, which consists of the *eBay*-like RUBiS benchmark [17] and a set of real world software, to compare FEMOSAA with its variants and other state-of-the-art approaches under the FIFA98 workload [18]. The results reveal its effectiveness and superiority over the others.

The contributions have clear impact on the synergy between software engineering for SAS and evolutionary computation as FEMOSAA combines the strengths form both fields. In particular, to control and exploit the power of EA/MOEA, software engineers only need to provide the feature model when using FEMOSAA, without being an expert on EA/MOEA. FEMOSAA also improves EA/MOEA and provides insights for EA/MOEA researchers to design better algorithms for SAS, since the identified elitist features and their dependencies serves as the systematic domain knowledge by which we can reduce the search space and guide the search, providing more chances for finding better solutions.

The reminder of the paper is organized as follows: Section II presents the background and the extended notions of numeric features. Section III gives an overview of FEMOSAA. Section IV illustrates our approach that transposes a feature model to EA/MOEA. Section V presents the MOEA/D-STM with dependency aware operators and knee selection. Experimental results, verifiability and threats to validity are discussed in Section VI. Finally, Sections VII and VIII present related work and conclusion respectively.

II. Background

A. Multi-Objective Evolutionary Algorithm

EA is a stochastic search-based meta-heuristic that has been widely used for multi-objective optimization, referred to as MOEA [19]. In EA/MOEA, the population contains a set of solutions (individuals), each of which is represented by a fixed-length thread-like chromosome carrying different values at each gene. During the search, by doing crossover and mutation on the solutions chosen from the mating selection, elite information is expected to propagate to their offspring. The survival selection ensures that the solutions having superior fitness, evaluated by the objective functions, can survive in the population for next iteration (generation). The ultimate goal of MOEA is to find a set of non-dominated solutions that approximates the *Pareto Front*. Notably, for every solution in the non-dominated set, any improvement of an objective will result in a degradation for at least one other objective. According to the strategy of survival selection, existing MOEAs can be divided into three categories: (i) Pareto-based method, e.g., NSGA-II [14]; (ii) indicator-based method, e.g., IBEA [20]; and (iii) decomposition-based method, e.g., MOEA/D [15]. FEMOSAA exploits the last method since it has shown better convergence and diversity than the prior two on many problems [15].

B. Feature Model with Numeric Features

The feature model [11], expressed as the tree structure, is a widely used notation for software engineers to represent the variability of a software [24]. Given its nature, there has been some successful attempts that apply feature model to design SAS [9] [10]. Figure 1 shows the feature model for the SAS we use in this work. As we can see there are four types of in-branch relation between a feature and its parent: $Optional$ refers to the feature might be deselected while $Mandatory$ denotes core features, which can not be deselceted. XOR represents the feature in a group such that exactly one group member can be selected; OR means a group that at least one group member needs to be selected. When a feature is selected, it means that such a feature is 'turned on'; similarly, deselection of a feature refers to it is 'turned off'. Selecting a feature implies that its parent should be selected too. In this work, we call a feature ***deselectable*** if it has $Optional$, OR or XOR relation to its parent; or ***conditionally deselectable*** if it has $Mandatory$ relation to its parent but there exist deselectable ancestors. There are also two cross-branch relations: (i) $F_i\ require\ F_j$ means the former can only be selected if the latter is selected; and (ii) $F_i\ exclude\ F_j$ denotes two features are symmetrically mutually exclusive. All those relations constitute to the dependency chain(s) in the model. The number of features in our example is 1151 with a search space of 1.3×10^{16}.

To better incorporate feature model with SAS and simplify the design, we distinguish ***categorical feature*** and ***numeric feature***. We define numeric features as: *a feature is numeric if it has more than one children in its XOR group and all its children can be quantified by number.* For example, in Figure 1, the *Memory* is clearly a numeric feature. Otherwise, the feature is categorical, e.g., *Cache Mode*. Similarly, a dependency is numeric as long as it is linked to numeric features and it involves quantitative comparisons. As in Figure 1, we propose the following cross-branch numeric dependencies for the engineers to specify in their design:

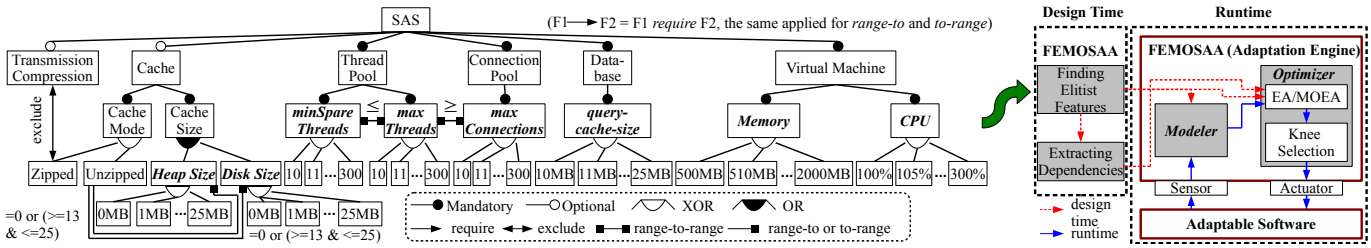

Fig. 1. The exampled feature model and the architecture of FEMOSAA. (the numeric features are shown in bold and italic letters)

- *range-to-range*. This is associated with two numeric features and it can be expressed as, e.g., F_i *range-to-range* F_j ($F_i < F_j$), meaning that F_i's selected child in its XOR needs to be smaller than that of F_j. It can be translated into categorical dependency easily: $F_i < F_j$ simply means F_i's XOR child C would have $exclude$ dependency on each F_j's XOR child, which is larger than or equal to C. Other quantitative comparisons (e.g., $>$) can be also applied.
- *to-range*. This constrains a categorical feature F_i w.r.t. a numeric feature F_j, e.g., F_i *to-range* F_j ($F_j < 10$), meaning that F_i can only be selected if F_j's selected child in its XOR falls in the given range, as expressed by mathematical formula. This can be translated to categorical dependency such that F_i would have $exclude$ dependency on each of F_j's XOR children that are not in the range.
- *range-to*. This is the inverse of *to-range* dependency.

Clearly, numeric dependencies can only be cross-branched while categorical ones exist on both in-branch and cross-branch. When a dependency is associated with one categorical feature and one numeric feature (i.e., *to-range* and *range-to*), we call it a hybrid dependency. Note that numeric features might have all types of dependencies but categorical features cannot be linked to numeric dependency.

III. FEMOSAA Overview

As shown in Figure 1, a SAS generally consists of two parts: an adaptable software that being managed at runtime, and an engine that controls the adaptation. The adaptable software could be a software stack that contains different inter-connected components or middleware, and its features form the feature model of the entire SAS. Our FEMOSAA framework is deployed as the adaptation engine and it operates on both design time and runtime of the SAS. At design time, FEMOSAA analyzes and transposes the feature model of SAS, which is provided by the software engineers, to the context of EA/MOEA. The outcome can be also used by the software engineers to implement related sensors/actuators, and to design the inputs of objective functions. FEMOSAA has two main components at runtime: a *Modeler* that contains the objectives functions, which can be updated on-the-fly using the data (i.e., objective values, utilization of features etc.) from sensors [22] [23]; and an *Optimizer* that realizes the EA/MOEA (i.e., MOEA/D-STM with knee selection), which is guided by the transposed information from feature model, to find a single optimized solution for adaptation via actuators. Given the uncertain and dynamic workload, these two components constitute the feedback loop that continually adapts the SAS towards better quality, i.e., *response time* and *energy consumption* in this work. The loop can be triggered either periodically or by thresholds on non-functional quality. Note that we consider the execution order of a solution as a separate issue from the optimization. Thus, given a valid and optimized solution, we assume that the valid order of execution, w.r.t. the dependency, is enforced in the actuators, which can be done through feature model analysis.

IV. Transposing Feature Model of Self-Adaptive Software to Evolutionary Algorithm

In this section, we present an automatic and systematic approach as part of FEMOSAA that transposes a feature model into EA/MOEA's context. At design time, the approach finds the necessary features (elitists) from the model to form chromosome representation and then extracts feature dependency w.r.t. the elitist features. Such information will be used at runtime to guide the evolutionary optimization.

To guarantee correctness of the transposition, it is imperative to ensure that the feature model has been fully tested and verified by existing tools [24]. Henceforth, this ensure that faults, e.g., dead features, false options and contradictory relations, have been already dealt with before the transposition. The verification of feature model is beyond the scope of this work, however. Unlike our work, the dependency related to numeric features is not treated explicitly in existing testing tools. However, as discussed in Section II-B, the numeric (and hybrid) dependency can be easily transferred into the categorical dependency, which can be then tested directly. We also assume that all possible children (including *0*) of numeric features are discretized and predefined.

A. Finding Elitist Features for Chromosome Representation

1) Growing the Feature Model Tree: To correctly transpose the feature model, we firstly grow the feature model tree for disclosing the hidden information inferred from deselectable features. This is achieved by adding children to a given feature F using the following steps in order:

- **G-1.** If F is a leaf feature that has OR relation to its parent, we then add two children representing *On* and *Off* in a XOR group to F.
- **G-2.** If F is a leaf feature that has neither OR nor XOR relation to its parent, we then add one child representing *On* in a XOR group to F.
- **G-3.** If F is a branch feature that has $Optional$, OR or XOR relation to its parent, we then add one child representing *Off* in a XOR group to F and to the descendants of F that are branch features.

After growing the tree, the added features and the steps that create them are shown in Figure 2.

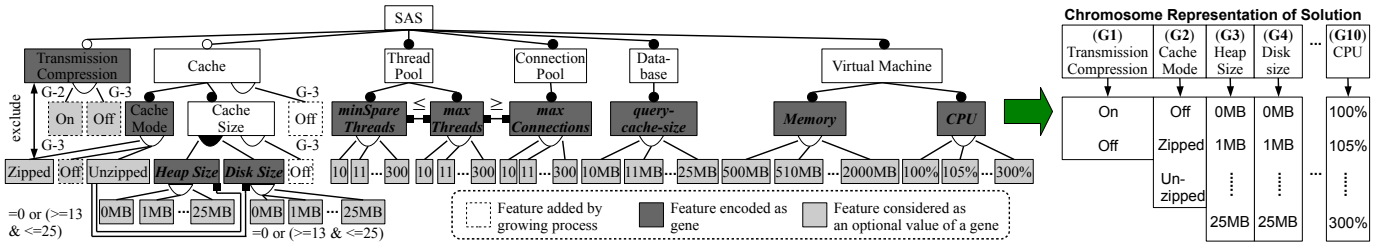

Fig. 2. The feature model after the process of finding elitist features and the resulted chromosome representation in form of genes (G_1, G_2, ..., G_{10}).

2) Identifying Genes from the Feature Model Tree: We have now obtained a model with no hidden information, the next phase is to find the elitist features for chromosome representation. Intuitively, we follow the rule that *a feature F is encoded as gene in the chromosome, if and only if, it is the parent of a XOR group, which contains more than one group member.* Subsequently, F's children within the XOR group constitute its set of optional values to be chosen in EA/MOEA, subject to the constraints in dependency. From now on, we call those features, which are chosen to be encoded in the chromosome, as ***genes***. It is easy to see that numeric features are always chosen as genes. As shown in Figure 2, there are 10 features being considered as genes. The others, which are not genes, can either have no implication on the variability, e.g., the *Thread Pool* feature is always mandatory; or they can be fixed to *On* in the SAS as their variability can be expressed by their descendants, e.g., the variability of *Cache Size* feature can be represented by the combination of *Heap Size* and *Disk Size* features. To make the model informative, we prune the features that are the only member in a XOR group. For all genes, if they select *Off* or *0* as their value then it means they are deselected; any other values mean they are selected. Note that when a gene selects *0*, it does not only mean its numeric value is set to 0, but also implies that the feature is 'turned off', which will have no further effects on the SAS.

B. Extracting Feature Dependency for Guiding Evolution

1) Analyzing and Refactoring the Dependency: Once we have identified the chromosome representation, we then need to extract dependency chain(s) to accommodate with the genes. Here, a single dependency between two genes represents the constraint on the dependent gene with respect to the conditions of main gene. The extracted dependencies and their imposed constraints are shown in Table I. We distinguish two categories of dependency: *in-branch* and *cross-branch*.

Extracting in-branch dependency chain(s) aims to handle the constraints introduced by $Optional$, $Mandatory$, OR and XOR relations within the chromosome representation. To achieve this, the features' in-branch dependencies are extracted in both vertical and horizontal direction.

Vertical analysis for extracting in-branch dependency helps to ensure that the in-branch relation between feature and parent are captured. For any feature F in the original feature model, we conduct the following vertical analysis:

- **VA-1.** If F is a gene and it is deselectable or conditionally deselectable, then for each path from F, the closest descendant gene D_g of F would have $require$ dependency on F, as D_g can not be selected without the presence of F. Additionally, if D_g has $Mandatory$ relation to its parent and the path between F and D_g does not contain deselectable features, then F would also have $require$ on D_g, as both features need to be selected at the same time.
- **VA-2.** In addition to **VA-1**, if F has XOR relation to its parent and it is a gene, then F would have a $require$ on its parent, denoted as $P_g = \alpha$ (F's parent P_g would always be a gene as ensured by our identifying process), where α is the reference of F in P_g; similarly, $P_g = \alpha$ would also have $require$ on F, as both features need to be selected at the same time. On the other hand, if F has XOR relation to its parent but it is not a gene, then for each path from F, the closest descendant gene D_g of F would have its own $require$ on $P_g = \alpha$. Under the same case, if D_g has $Mandatory$ relation to its parent and the path between F and D_g does not contain deselectable features, then $P_g = \alpha$ would also need to have $require$ on D_g.
- **VA-3.** In addition to **VA-1**, if F has OR relation to its parent, we find the closest deselectable ancestor of F, denoted as A, if such a A does exist. Now, if there exist at least one ancestor of F, which is gene; or F's parent is neither deselectable nor conditionally deselectable; or there exist at least one closest descendant gene, D_g, of a path from A, such that D_g has $Mandatory$ relation to its parent and there is no deselectable features in the path between D_g and A, then this means that the OR group, for which F's parent is the root, needs to select at least one group member. Thus, unless there already exist a $at\text{-}least\text{-}one\text{-}exist$ dependency, F (if F is a gene) or its closest descendant genes, each of which follows different paths (if F is not a gene), would have $at\text{-}least\text{-}one\text{-}exist$ on (i) the other closest descendant genes of F if it is not gene; (ii) those sibling genes of F in the same OR group; and (iii) the closest descendant genes, each of which follows different paths from the F's siblings that are not genes but in the same OR group as F.

The horizontal refactoring, on the other hand, is to ensure the elimination of deselectable features does not affect the original variability. Suppose that F is a feature in the original feature model and that A is the closest deselectable ancestor of F, if such a A does exist. Now, assuming that A is not a gene, and that there is no gene on the path from A to F, we then conduct the following horizontal refactoring:

- **HR-1.** If F has $Optional$ relation to its parent then we do nothing, even if it is a gene. This is because the selection of F does not affect other descendant genes of A.
- **HR-2.** If F is a gene that has $Mandatory$ relation to its parent and it is conditionally deselectable, then for each path from A, the closest descendant gene D_g of A (excluding F

TABLE I. THE EXTRACTED DEPENDENCY CONSTRAINTS BETWEEN TWO GENES AND THE RELATED SET OPERATIONS FOR MERGING DEPENDENCY.

Dependency (denoted as D)	Constraints on G_i	Merging with Another Dependency D'
G_i *require* G_j	if $G_j = Off or 0$, then $A_s = \{Off\} or \{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha$ *require* G_j	if $G_j = Off$ *or* 0, then $A_s = A - \{\alpha\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i *require* $G_j = \alpha$	if $G_j \neq \alpha$, then $A_s = \{Off\}$ *or* $\{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha_1$ *require* $G_j = \alpha_2$	if $G_j \neq \alpha_2$, then $A_s = A - \{\alpha_1\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i *exclude* G_j	if $G_j \neq Off or 0$, then $A_s = \{Off\} or \{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha$ *exclude* G_j	if $G_j \neq Off$ *or* 0, then $A_s = A - \{\alpha\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha_1$ *exclude* $G_j = \alpha_2$	if $G_j = \alpha_2$, then $A_s = A - \{\alpha_1\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i *at-least-one-require* G_j	if $G_j = Off or 0$, then $A_s = \{Off\} or \{0\}$. Otherwise, $A_s = A$.	if D' is (or merged from only) *at-least-one-require* that related to the same root of OR group as D's, then $A_s \cup A'_s$. Otherwise, $A_s \cap A'_s$.
$G_i = \alpha$ *at-least-one-require* G_j	if $G_j = Off$ *or* 0, then $A_s = A - \{\alpha\}$. Otherwise, $A_s = A$.	if D' is (or merged from only) *at-least-one-require* that related to the same root of OR group as D's, then $A_s \cup A'_s$. Otherwise, $A_s \cap A'_s$.
G_i *at-least-one-exist* G_j	if $G_j = Off or 0$, then $A_s = A - \{Off\}$ *or* $A - \{0\}$. Otherwise, $A_s = A$.	if D' is (or merged from only) *at-least-one-exist* that related to the same root of OR group as D's, then $A_s \cup A'_s$. Otherwise, $A_s \cap A'_s$.
G_i *range-to-range* G_j (e.g., $G_i < G_j$)	if $G_j = \alpha$, then $A_s = \{\alpha_1, ..., \alpha_n\}$ where $\forall \alpha_n \in A$, and $\forall \alpha_n$ meets the given condition w.r.t. α, e.g., $\forall \alpha_n < \alpha$, etc.	$A_s \cap A'_s$.
G_i (R) *range-to* G_j	if $G_j = Off$ *or* 0, then $A_s = A - R$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i (R) *range-to* $G_j = \alpha$	if $G_j \neq \alpha$, then $A_s = A - R$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i *to-range* G_j (R)	if $G_j = \alpha$; $\alpha \notin R$, then $A_s = \{Off\} or \{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha_1$ *to-range* G_j (R)	if $G_j = \alpha_2$; $\alpha_2 \notin R$, then $A_s = A - \{\alpha_1\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
Additionally, whenever the set operations lead to empty set, we fix $A_s = \{Off\}$ *or* $\{0\}$.		

G_i and G_j are dependent and main gene, respectively; A is the entire set of optional values for G_i; A_s denotes the set of values for G_i, given a selected value of G_j; α, α_1, α_2 and α_n denote some selected values for G_i or G_j; R is a given constrained set of range, e.g., $G_j < 10$ etc; D' is another single or merged dependency for which G_i is the dependent gene; A'_s denotes the set of values for G_i, given the selected value(s) of the main gene(s) in D'.

itself) would have $require$ on F. This can ensure whenever D_g is selected, F would be also selected.

- **HR-3.** If F has OR relation to its parent and it does not have $at\text{-}least\text{-}one\text{-}exist$ dependency for the group, then for each path from A (except the paths that pass through F's OR group), the closest descendant gene D_g of A would have $at\text{-}least\text{-}one\text{-}require$ on F, if F is a gene; or on those closest descendant genes of F, each of which follows different paths from F, if F is not a gene. Hence, whenever D_g is selected, at least one member of the OR group of F (or their closest descendant genes) will be also selected.
- **HR-4.** If F has XOR relation to its parent then we do nothing, even if it is a gene. This is because our identifying process ensures that the parent of F would always be a gene, which also express the selection of F.

After considering the in-branch dependency, we now focus on refactoring the cross-branch dependency. If both sides of a cross-branch dependency are genes, then it can be extracted directly. However, if either side (or both) of the feature is not a gene, then treatment is needed. Suppose that a feature F is associated with one or more cross-branch dependency and that F is not a gene, we then do the following refactoring:

- **CR-1.** If F is a branch, then its cross-branch dependencies are migrated to those closest descendant genes of F, each of which follows different paths from F. Further, if F is the root of a OR group and it is the main feature in any $require$ dependency, then those $require$ would be changed to $at\text{-}least\text{-}one\text{-}require$, which are migrated to the member genes of F's OR group, and to the closest descendant genes, each of which follows different paths from those members that are not genes.
- **CR-2.** If F is a leaf, then its cross-branch dependencies are migrated to the parent of F, denoted as P_g, where the dependency would remain the same but the main gene becomes $P_g = \alpha$, where α is the reference of F in P_g. Here, P_g is always a gene, because it would be a design error, if F was to has $Mandatory$ relation to P_g. In addition, our growing process has ensured that F has neither OR nor $Optional$ relation to P_g while it is a leaf.

Finally, the extraction which occurred on the model and the extracted dependency chain(s), w.r.t. the elitist chromosome, are shown in Figure 3. The constraint of a dependent gene imposed by a dependency, according to Table I, can be expressed using a ***value tree***, where each leaf is a set of optional values constrained by the corresponding condition in a branch, i.e., the selected values of main gene. In Figure 3 (right), the value tree for the dependency between *Cache Mode (G_2)* gene and *Transmission Compression (G_1)* gene constrains that the former can only be *Off* or *Unzipped* if the latter is *On*; or any optional values, otherwise. Note that if a gene is not a dependent of any dependency, then it would have a value tree without branches.

2) Merging the Dependency: After the extraction, we can see that a dependent gene might have multiple dependencies on the same or different main genes. To construct a combined value tree for a dependent gene, the dependencies, by which it is constrained, need to be merged one by one using set operators (union or intersection) to combine the leaves from their value trees. Table I shows what set operators are needed for each dependency type when merging with the others. To ensure correctness of merging, the dependencies that require union are always merged ahead of the others, leading to a set of partially merged dependencies. Since each of them are merged from identical dependency type, as shown in Table I, they can then be merged with each others (and the remaining single dependencies) using the same set operators as if they are single dependencies.

Specifically, for a given dependent gene, if it has two or more dependencies with identical main genes, then the leaves, which are constrained by the same condition in the

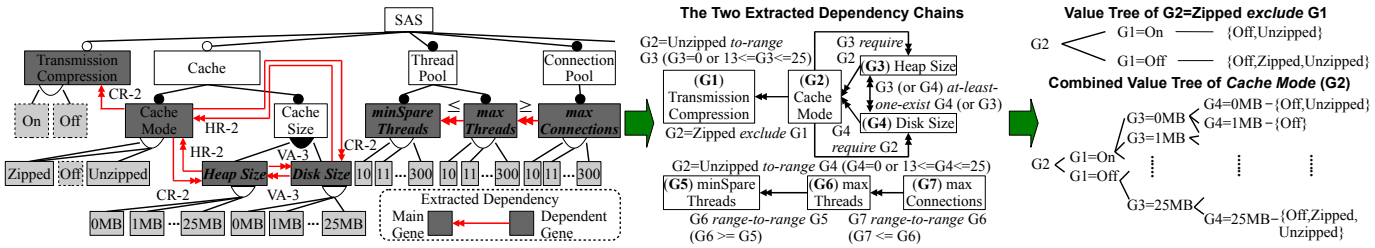

Fig. 3. The extraction of dependency chains for genes, example of value tree and dependency merging.

branches, would be combined directly. Afterwords, for those dependencies that have different main genes, all branch nodes of a single or already combined value tree are replicated and grafted (as a whole) to each right-most branch node of another single or already combined value tree, forming new levels for the newly combined value tree, representing the combinatorial conditions. Then, for the two value trees that were grafted, their leaves, whose original ancestors are now on the same path from root to a right-most node in the newly grafted tree, are combined to create the new leaf-set using the set operations shown in Table I. The process stops when all related dependencies are merged and their value trees are combined, resulting a finally combined value tree. Figure 3 (right) illustrates the finally combined value tree for *Cache Mode* (G_2) gene, which originally contains dependencies on three different main genes.

It is worth noting that, despite the given feature model is contradictions free, it might still possible for the combined leaves to have empty set under some conditions. For example, recall *Disk Size* gene from Figure 3, when its $require$ dependency on *Cache Mode* merges with its $at\text{-}least\text{-}one\text{-}exist$ on *Heap Size*, then combing the leaves for the condition where *Cache Mode* selects *Off* and *Heap Size* selects *0 MB* would lead to an empty set. This is due to the fact that some dependencies have different priorities: in the above example, $at\text{-}least\text{-}one\text{-}exist$ would constrain *Disk Size* only if it is allowed to be selected as indicated by $require$. In those cases, we fix $\{Off\}$ (or $\{0\}$) as the new leaf-set for the dependent gene, representing the fact that it needs to be 'turned off'and cannot affect the SAS under the coressponding conditions.

In Section V, we will describe how the elitist chromosome, the extracted dependency chain(s) and the (combined) value trees of genes are seamlessly injected into the EA/MOEA, regardless to the order of changes on genes.

V. Feature Guided MOEA/D-STM with Knee Selection

Unlike the classic aggregation methods which only specify and solve one aggregation at a time, MOEA/D decomposes the original multi-objective problem into several subproblems, each of which is an additive or Tchebycheff [25] aggregation with automatically generated weights. Then, it leverages the population-based technique in classic EA/MOEA to solve these subproblems collaboratively. FEMOSAA exploits the recent MOEA/D variant, MOEA/D-STM, which achieves better balance between convergence and diversity by modifying the survival selection of the original MOEA/D. In a nutshell, MOEA/D-STM treats subproblems and solutions as two sets of agents. Each agent has its preferences over agents on the other side. In particular, subproblems concern convergence while solutions concern diversity. The selection result is the stable matching between subproblems and solutions, which achieves an equilibrium between their preferences. In the following, we briefly explain the main procedures of MOEA/D-STM, while interested readers can refer to [16] for more details.

- **Step 1.** Initialize the population and numbers of evenly distributed weight vectors [26], each of which represents a subproblem (we use Tchebycheff [25]), and build their fixed size of neighborhood structure using Euclidean distance.
- **Step 2.** Generate offspring via crossover and mutation among parents randomly selected from neighborhoods.
- **Step 3.** Combine the parents and offspring together and build a stable matching relationship between subproblems and solutions. Solutions, which have been successfully matched with the subproblems, are selected to survive. Repeat from **Step 2** if stopping criterion is not met.

In FEMOSAA, we integrate the extracted dependencies and the (combined) value trees of genes into **Step 2**; and apply the knee selection proposed in Section V-D for the posterior decision making following **Step 3**.

A. Objective Functions

The elitist chromosome representation from Section IV forms the structure of solutions in EA/MOEA, it also helps to define the inputs of the objective (fitness) functions used in the optimization. Formally, at a given timestep t, the objective functions that we consider are:

$$RT(t+1) = f(G_1(t+1), G_2(t+1), ..., G_n(t+1), \delta(t)) \quad (1)$$

$$EC(t+1) = g(G_1(t+1), G_2(t+1), ..., G_n(t+1), \delta(t)) \quad (2)$$

whereby $G_1(t+1), G_2(t+1), ..., G_n(t+1)$ denote the genes and their selected values in a possible adaptation solution; $\delta(t)$ denotes a set of most recent environmental factors, e.g, workload in this work. Given an adaptation solution, $RT(t+1)$ and $EC(t+1)$ are the expected fitness values for *response time* and *energy consumption* respectively.

In this work, we adopt the machine learning based model from [22] [23] to build the objective functions, since it relies on few assumptions of the application domain. However, other modeling approaches [27] [28] can be easily applied.

B. Dependency Aware Mutation Operator

At runtime, we use a modified bit-flip mutation operator as the basis, in which each gene can be randomly mutated to one of its optional values subject to a mutation rate. As illustrated

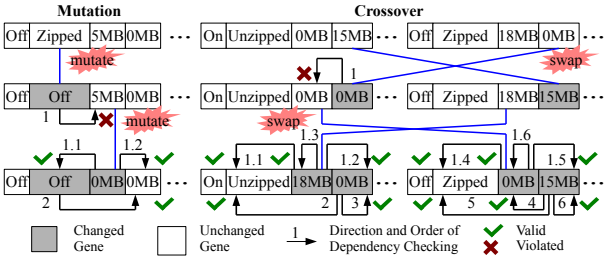

Fig. 4. Workflow of dependency aware mutation and crossover operators. (The exampled genes from left to right represent *Transmission Compression*, *Cache Mode*, *Heap Size* and *Disk Size*)

in Figure 4, our extended dependency aware mutation operator has the following recursive steps:

- **Step 1.** When a gene G (*Cache Mode*) needs to be mutated, we randomly select a value from its (combined) value tree w.r.t. the selected values of G's main genes (*Transmission Compression*, *Heap Size* and *Disk Size*).
- **Step 2.** Then, we propagate, according to the dependency chain, to G's dependent genes (*Heap Size* and *Disk Size*) and we then validate if those genes violate any dependency using their (combined) value trees. If violation found, we would mutate the corresponding gene and start from **Step 1** for it (e.g., *Heap Size=5MB* in Figure 4).
- **Step 3.** The process stops when no violation found.

Using this operator, we guarantee that the mutation process in EA/MOEA would be better guided, such that only valid adaptation solutions can be explored regardless to the order of mutation. Since the given feature model has been fully tested, there will be at least one state such that all genes can satisfy all dependencies, which in turn, preventing infinite loop in the presence of circular dependencies.

C. Dependency Aware Crossover Operator

As shown in Figure 4, we use uniform crossover where two genes, each of which from a different parent and both are at the same position in the chromosome, can be swapped subject to a crossover rate. After injection, our dependency aware crossover operator uses the following recursive steps:

- **Step 1.** When a gene G (*Disk Size*) is swapped in an offspring, we propagate, according to the dependency chain, to G's dependent genes (*Heap Size* and *Cache Mode*). If the dependent gene has not been swapped and it violates the dependency, we then swap the gene and repeat **Step 1** for it (e.g., *Heap Size=0MB* in Figure 4).
- **Step 2.** Next, we check if G itself violates dependencies and if any of its main genes have not been swapped. When both conditions are true, we then swap those main genes (*Cache Mode* and *Heap Size*, but this is not needed in the example shown) and start from **Step 1** for each of them.
- **Step 3.** The process terminates when there is no dependency violation or all genes in the parents are swapped.

In this way, we guarantee that only valid adaptation solutions would be produced, given any order of the crossover.

D. Knee Selection

To achieve a more balanced trade-off when optimizing SAS without human intervention, we have developed a knee selection method for the posterior decision making. As shown

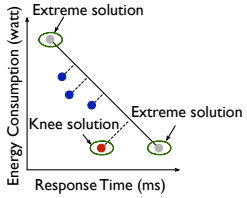

Fig. 5. Knee solution(s) in the non-dominated set.

in Figure 5, in the final non-dominated set produced by feature guided MOEA/D-STM, we firstly construct a line ℓ that connects the extreme solutions holding the worst value at each single objective. Then, we calculate the perpendicular distance from each non-dominated solution $\mathbf{x}$ to the ℓ:

$$d(\mathbf{x},\ell) = \begin{cases} \frac{\epsilon}{\sqrt{a^2+b^2}}, & if \quad \epsilon < 0 \\ -\frac{\epsilon}{\sqrt{a^2+b^2}}, & otherwise \end{cases} \tag{3}$$

where $\epsilon = a \times f(\mathbf{x}) + b \times g(\mathbf{x}) + c$ and the parameters a, b and c can be identified by the extreme solutions. The solution(s), which has the largest perpendicular distance to ℓ, is the knee solution(s) that we are seeking. When there are multiple knee solutions, we randomly select one for adaptation. Note that, by replacing ℓ with a surface, we can easily extend our knee selection method for more than two objectives.

VI. Experiments and Evaluation

We have conducted comprehensive experiments to evaluate the effectiveness of FEMOSAA by means of comparing it with its variants and the state-of-the-art under different metrics. Our evaluation aims to answer the following research questions:

- **RQ1.** What is the added value of the elitist chromosome representation in contrast to the binary representation?
- **RQ2.** What are the benefits that the dependency aware operators and the knee selection mechanism provide?
- **RQ3.** What is the effectiveness of FEMOSAA in contrast to other state-of-the-art search-based frameworks?
- **RQ4.** Does MOEA/D-STM provides comparative results when compared with NSGA-II, single objective Genetic Algorithm (GA) and the exact algorithm?
- **RQ5.** What is the overhead of FEMOSAA, in terms of execution time, compared to other state-of-the-art approaches? Which part(s) cause the most overhead?

A. Verifiability and Methodology

We conducted experiments on a dedicated server, which runs Ubuntu Linux 14.04 and has Intel i5 2.8GHz Quad Cores, 4GB RAM. To achieve separation between adaptation engine and the adaptable software, we use Xen v3.0.3 [29] as the hypervisor to create a virtulaized environment. We have implemented FEMOSAA as the adaptation engine using Java, JDK 1.6, and it is deployed on the *Dom0* of Xen. For setting up MOEA/D-STM, we use a population size of 100 for 10 generations; the mutation and crossover rate are 0.1 and 0.9 respectively. The number of evenly distributed weight vectors is 100 and the size of neighborhood is 20. Those settings are either common values or have been tailored for runtime optimization, w.r.t. quality and overhead. To mitigate interference, we use one vCPU and 800MB RAM on *Dom0*.

TABLE II. COMPARING FEMOSAA WITH ITS VARIANTS. (THE BEST RESULT AND STATISTICAL SIGNIFICANCE ARE HIGHLIGHTED. GM DENOTES GEOMETRIC MEAN; ES MEANS EFFECT SIZE; HV AND ED ARE THE MODIFIED HYPERVOLUME AND EUCLIDEAN DISTANCE RESPECTIVELY.)

FEMOSAA and Variants	Response Time (ms)			Energy Consumption (watt)			Valid Soltions (%)	HV	ED
	GM	p value(ES) with FEMOSAA	p value(ES) with FEMOSAA-KD	GM	p value(ES) with FEMOSAA	p value(ES) with FEMOSAA-KD			
FEMOSAA	**38.39**	N/A	***p*=0.030(medium)**	**3.13**	N/A	***p*=0.002(medium)**	**100**	**1**	**0**
FEMOSAA-D	115.59	***p*=0.002(medium)**	p=0.384(trival)	3.60	p=0.153(small)	p=0.516(trival)	13.29	0.30	0.33
FEMOSAA-K	103.99	***p*=0.003(medium)**	p=0.638(trival)	3.26	p=0.535(trival)	p=0.051(small)	**100**	0.46	0.26
FEMOSAA-KD	83.62	***p*=0.030(medium)**	N/A	3.80	***p*=0.002(medium)**	N/A	9.39	0.41	0.26
FEMOSAA-0/1	168.47	***p*=0.000(large)**	***p*=0.042(small)**	4.91	***p*=0.000(large)**	***p*=0.018(medium)**	0	0	0.71

The adaptable software is a software stack that contains RUBiS [17], which is a well-known software benchmark that simulates the *eBay* model, and a set of real world software including Tomcat v6.0.28 [3], MySQL v3.23.58 [30] and Linux kernel v2.6.16.29 running on a configurable guest virtual machine. Ehcache v2.6 [31] is plugged to RUBiS as the cache management module. All the features that can be adapted at runtime are represented by the feature model in Figure 1. To simulate a realistic workload within the capacity of our testbed, we vary the number of clients according to the compressed FIFA98 workload [18], as shown in Figure 6. This setup can generate up to 600 parallel requests, which are sufficiently significant. The workload is generated by another machine using the client emulator provided in RUBiS. The source code of FEMOSAA and the adaptable software can be publicly accessed via GitHub [32] [33].

As discussed, FEMOSAA finds elitist features and extracts their dependency at design time, then the outcomes are passed to MOEA/D-STM and knee selection for runtime optimization. The search space, after reduced by FEMOSAA, is still around 1.2×10^{15}. To stabilize the training of objective functions, we firstly run the SAS using randomly generated workload for 150 intervals. Next, the trained objective functions are used and updated at runtime under the FIFA98 trace for 102 intervals. For all experiments, the sampling interval is 120s and we trigger an optimization run by the end of each interval.

B. Effectiveness

1) Comparing with Variants of FEMOSAA: To answer *RQ1* and *RQ2*, we compare FEMOSAA with its variants:

FEMOSAA-D—This is similar to our FEMOSAA except that it does not use dependency aware operators. Therefore, if the final population contains invalid solutions, it automatically filters them and work on the valid ones only. When no valid solutions found in final non-dominated set, it corrects the invalid solutions via dependency aware mutation operator.

FEMOSAA-K—Another variant of FEMOSAA without knee selection. Hence, in the final non-dominated set, one solution is randomly selected for adaptation.

FEMOSAA-KD—This variant neither consider dependency nor knee selection in the optimization. Hence, it uses the ad hoc strategies for FEMOSAA-D and FEMOSAA-K.

FEMOSAA-0/1—A baseline variant, built from FEMOSAA-KD, that exploits the classic binary chromosome representation using all the features from Figure 1. It neither considers dependency nor knee in the optimization. Note that if more than one members in a XOR group are selected, we then randomly choose one to create a valid solution.

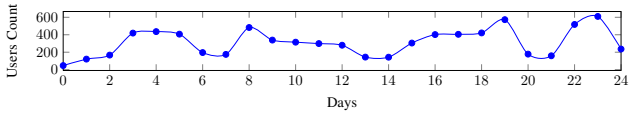

Fig. 6. The compressed FIFA98 workload (from June 7 to July 1, 1998).

For each interval, we record the average *response time* for all requests to the SAS and we use PowerAPI [34] for measuring the average *energy consumption* on the SAS. We report on the Geometric Mean (GM) for all the observed objective values over 102 intervals. We also compare the average valid adaptation solutions found over all the intervals. Finally, to compare the overall quality for both objectives, we apply a modified Hypervolume (HV) [21] and Euclidean Distance (ED) [35] on GMs to assess the balance in trade-off and the extent to which both objectives are optimized, respectively. The GMs are normalized before computing those two metrics so that they range from 0 to 1. Hence, the modified HV is computed as $\prod_{i=1}^{n}(1-\frac{GM_i-GM_{i,b}}{GM_{i,w}-GM_{i,b}})$ and ED can be found by $\frac{1}{n}\times\sqrt{\sum_{i=1}^{n}(\frac{GM_i-GM_{i,b}}{GM_{i,w}-GM_{i,b}}-0)^2}$, where n is the number of objectives and GM_i is the GM for the ith objective; $GM_{i,b}$ and $GM_{i,w}$ are the best and worst GMs that we observed for the ith objective, respectively.

To confirm statistical significance of the comparisons on *response time* and *energy consumption*, we performed *Wilcoxon Signed-Rank* test (two-tailed) for any comparison between FEMOSAA (or FEMOSAA-KD) and the others, since our data does not follow Gaussian distrbution. We use 95% as the confidence interval ($\alpha = 0.05$) and this means that, if the test produces a p that is smaller than 0.05, then we can reject the null hypothesis H_0, which states that the given two approaches cannot be statistically distinguished when optimizing *response time* or *energy consumption* of the SAS. The *Effect Size (ES)* for each test is also reported and we follow the categories in [36] to measure the meaningfulness of effect size.

As shown in Table II, FEMOSAA achieves the best results for all metrics. In particular, FEMOSAA and FEMOSAA-KD outperform FEMOSAA-0/1 on both *response time* and *energy consumption*, with $p < 0.05$ and at least small effect size. They also obtain better HV and ED. These results confirm the benefits of our elitist chromosome representation. We can also see that FEMOSAA is better than FEMOSAA-D on both objectives, together with better HV and ED values, but the improvement on *energy consumption* is barely significant. Similar observation can be registered when comparing FEMOSAA with FEMOSAA-K. In addition, FEMOSAA yields statistically better results than those of FEMOSAA-KD on all metrics with medium effect size. All those facts revel the benefits of dependency aware operators and knee selection in the optimization for preventing unneeded function evaluations on

TABLE III. COMPARING FEMOSAA WITH THE STATE-OF-THE-ART. (THE BEST RESULT AND STATISTICAL SIGNIFICANCE ARE HIGHLIGHTED. GM DENOTES GEOMETRIC MEAN; ES MEANS EFFECT SIZE; HV AND ED ARE THE MODIFIED HYPERVOLUME AND EUCLIDEAN DISTANCE RESPECTIVELY.)

Search-based Optimization Approach	Response Time (ms)			Energy Consumption (watt)			Valid Solutions (%)	HV	ED
	GM	p value(ES) with FEMOSAA	p value(ES) with FEMOSAA-KD	GM	p value(ES) with FEMOSAA	p value(ES) with FEMOSAA-KD			
FEMOSAA	**38.39**	N/A	***p*=0.030(medium)**	**3.13**	N/A	***p*=0.002(medium)**	**100**	**1**	**0**
FEMOSAA-KD	83.62	***p*=0.030(medium)**	N/A	3.80	***p*=0.002(medium)**	N/A	9.39	0.80	0.08
DUSE	118.57	***p*=0.004(medium)**	p=0.529(trival)	4.81	***p*=0.000(large)**	***p*=0.020(medium)**	10.25	0.63	0.15
PLATO	157.48	***p*=0.000(large)**	***p*=0.043(small)**	4.88	***p*=0.000(large)**	***p*=0.012(medium)**	8.39	0.52	0.21
FUSION	336.79	***p*=0.000(large)**	***p*=0.000(large)**	15.81	***p*=0.000(large)**	***p*=0.000(large)**	64.71	0	0.71

invalid solutions and searching for balanced trade-off, respectively. Interestingly, when comparing FEMOSAA-KD with FEMOSAA-D and FEMOSAA-K, the first has better result on *response time* but worse on *energy consumption*. However, their differences are not statistically significant and the HV and ED values are also closed. Such result implies that using dependency aware operators or knee selection independently might not provide statistically significant benefits. Finally, we can see that our dependency aware operators guarantee only valid solutions are explored, while in the absence of those operators, only less than 15% are valid, which could negatively affect the quality of optimization.

For *RQ1*, we can conclude that the elitist chromosome representation helps to provide statistically better optimization results, it can also reduce overhead, as we will show in Section VI-C. On the other hand, better quality and more balanced trade-off can be obtained by explicitly considering both dependency and knee in EA/MOEA, which answers *RQ2*.

2) Comparing with State-of-the-art: To answer *RQ3* and *RQ4*, we compare FEMOSAA and FEMOSAA-KD with the following state-of-the-art:

DUSE [5]—A representative framework that optimizes SAS using NSGA-II. Since it does not consider dependency and knee in the optimization, we adapt the ad hoc strategies applied by FEMOSAA-D and FEMOSAA-K.

PLATO [4]—An approach that applies weighted sum of objectives and GA in the optimization. It also does not consider dependency in the optimization and thus we use the same ad hoc strategy as FEMOSAA-D. We specify equal weights on the objectives to find the single best solution.

FUSION [9]—A well-known feature model based framework that formulates the SAS optimization as an Integer Programming problem, which assumes aggregate objective function and is resolved by an exact algorithm (we use *Branch-and-Bound* in the experiments). FUSION considers categorical dependency only, thus when no valid solution found, we then fix the violations on numeric dependency. To avoid unacceptable execution time, we forcibly return the best solution so far when it hits a time threshold, i.e., 40s.

Note that we have focused on implementing the optimization algorithms designed in those state-of-the-art approaches, the other parts, e.g., the modeling component and the parameters of EA, are assumed to be identical to FEMOSAA in all experiments. To ensure fair comparison, the elitist chromosome representation is used in both DUSE and PLATO. We apply the same set of metrics, statistical test and method for categorizing effect size as discussed eariler.

From Table III, we can see that FEMOSAA again achieves the best results on all metrics. Particularly, it outperforms the state-of-the-art approaches on both objectives with $p < 0.05$ and at least medium effect size. We note that it yields significantly better results than PLATO, which is constrained by its nature of weighted sum objective function; and FUSION, which is incapable to optimize efficiently under large search space. These confirm the superiority of FEMOSAA over the others. Similar observation can be obtained when comparing FEMOSAA-KD with the state-of-the-art approaches, though its improvement on *response time* against DUSE (with NSGA-II) is not statistically significant. This confirms the superiority of multi-objective optimization and that in contrast to NSGA-II, MOEA/D-STM can better explore good solutions that might be spread over the search space. Finally, since the state-of-the-art approaches do not consider all dependency, their ratio of valid solutions ranges from 8.39% to 64.71%.

For *RQ3* and *RQ4*, we can conclude that FEMOSAA yields better optimization results than the state-of-the-art and that MOEA/D-STM helps to efficiently find better solutions than NSGA-II, GA and exact algorithm.

C. Overhead

To answer *RQ5*, we compare FEMOSAA with FEMOSAA-KD, FEMOSAA-0/1 and state-of-the-art approaches in terms of the average execution time. As shown in Figure 7, FEMOSAA (1.96s) has similar execution time to FEMOSAA-KD (2.04s), meaning that the dependency aware operators and knee selection generate negligible overhead. Additionally, FEMOSAA has significantly smaller execution time than that of FEMOSAA-0/1, as the latter encodes a high dimension of genes in chromosome. This is another added value of the elitist chromosome representation, which is an additional answer to *RQ1*. On the other hand, FEMOSAA produces somewhat bigger overhead than DUSE and PLATO. We believe that this is caused by MOEA/D-STM where the stable matching mechanism poses extra complexity, especially when the solutions found have diverse quality. Finally, FUSION has the heighest overhead as exact algorithm fails to scale with large search space. Therefore, to answer *RQ5*, we can conclude that FEMOSAA has a lower overhead than FUSION, but higher than DUSE and PLATO; this is due to the complexity of stable matching when balancing convergence and diversity.

D. Threats to Validity

Construct validity arises from the bias introduced by the stochastic nature of the considered EAs/MOEAs in experiments, especially for those that do not account for dependency and knee in the optimization. Indeed, to draw meaningful conclusion, repeated runs are required for stochastic search-based

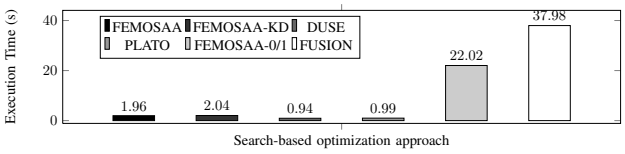

Fig. 7. The overhead comparison.

optimization. We have mitigated such bias by conducting 102 optimization runs for FEMOSAA and the others following the FIFA98 workload trace. We have also reported statistical test results and effect sizes of these repeated runs.

Internal validity is related to the values of parameters for MOEA/D-STM and other EAs/MOEAs. The setup in this work has been carefully tailored such that it produces good trade-off between the quality of optimization and the overhead. However, these values might vary depending on the characteristics in the context, e.g., the given feature model, the types of SAS and the environmental conditions.

External validity is linked to the benchmark and scenario that are used in the experiments. Although we have diversified the adaptable software by introducing a stack of software, it may still represent limited characteristic, i.e., web-based software that is I/O intensive. However, building the experiments and deploying the SAS are non-trivial, which could involve potential expenses. In future work, more experiments on other SASs will be conducted, e.g., CPU intensive software.

VII. Related Work

Search-based optimization has been widely applied for SAS, either constitutes a general framework or being tailored to a specific domain, e.g., service systems or cloud systems. A common way of optimizing multiple objectives is simple aggregation (e.g., weighted sum) such that they can be resolved in a single objective function. Hence, exact algorithms are often utilized for SAS [9] [37] [13], as it guarantees to find the optimal solution. FUSION [9] is a framework that optimizes SAS using exact algorithm. At design time, it also applies feature model to represent the software. However, unlike FEMOSSA, it does not consider numeric features and they formulate the problem as Integer Optimization problem, which is solved by binary encoding of all the features. [37] also follow similar idea of FUSION but they additionally use Fuzzy logic to constitute the objective functions. MOSES [13] is a framework that is designed for self-adaptive service systems, where the optimization is formulated as Linear Programming problem that is convex and can be solved exactly. The major issue of exact algorithm is that they fails to scale when the search space is large. Further, they often work on an objective aggregation, which limits their applicability and capability.

There exists some work that leverages single objective EAs for optimizing SAS based on an objective aggregation. PLATO [4] and VALKYRIE [12] are two examples. However, as opposed to the automatic approach in FEMOSSA, they have relied on manual encoding of the SAS into chromosome representation and they do not consider dependency between features. Further, it is well-known that the relative weights are difficult to be tailored by the engineers; and a single aggregation could restrict the search, causing limitation when searching for good solutions spread over the search space.

Though exploiting MOEAs is an emerging trend for SAS, our use of MOEA/D-STM and knee selection represent new attempts for SAS optimization at runtime. Among others, DUSE [5] uses NSGA-II to produce a set of non-dominated solutions for SAS. Similarly, NSGA-II has also been applied in [38] [6] [7], which focuses on cloud systems that adapt resource provision at runtime. Nevertheless, their encoding of the SAS has been manual and no dependency between features are considered. Furthermore, they rely on the non-dominated set, from which any solution can be used for adaptation, which can entail imbalanced trade-off. FEMOSSA, on the other hand, exploits MOEA/D-STM, which combines the benefits of mutiple objectives aggregation and solutions' population. FEMOSSA also find a knee solution that is generally preferable. Further, FEMOSSA relies on automatic approach where the elitist features are identified and the dependency is extracted to guide MOEA/D-STM. Pascual et al. [10] proposed a framework for self-adaptive mobile system, which uses feature model with MOEA and feature dependencies are considered in the optimization. However, it relies on binary representation of all features and only categorical features are considered. Further, they proposed to repair the invalid solutions found by the mutation and crossover. On the contrary, our dependency injection to those operators prevents invalid solutions from being explored. They do not consider knee selection either.

VIII. Conclusion

This paper presents FEMOSAA, a novel framework that automatically synergizes feature model and EA/MOEA, to optimize SAS at runtime. At design time, FEMOSAA finds elitist features, including categorical and numeric ones, and extracts their dependencies, which are injected into MOEA/D-STM for runtime optimization. The feature model serves as the domain knowledge that can reduce the search space and guide the search, henceforth increasing the chance for finding better solutions. Further, FEMOSAA finds knee solutions that achieve a balanced trade-off. By comparing FEMOSAA with its variants and the state-of-the-art on a real world SAS, we show that FEMOSAA produces statistically better and more balanced results for trade-off with reasonable overhead.

Our work impacts and advances the synergy between software engineering for SAS and evolutionary computation: with FEMOSAA, software engineers can exploit EAs/MOEAs to tackle SAS optimization without prior extensive expertise of the EA/MOEA. On the other hand, automatic transposition of feature model into EA/MOEA's context can improve EA/MOEA, and make the domain knowledge systematic and comprehensible for EA/MOEA researchers, who in turn can design more effective algorithms for SAS. In future work, we plan to apply FEMOSAA in other domains of SAS and extend it for managing more conflicting objectives.

Acknowledgment

The research was supported by the DAASE Programme Grant from EPSRC (EP/J017515/1).

REFERENCES

[1] R. de Lemos, H. Giese, H. A. Müller, M. Shaw, J. Andersson, M. Litoiu, B. Schmerl, G. Tamura, N. M. Villegas, T. Vogel, D. Weyns, L. Baresi, B. Becker, N. Bencomo, Y. Brun, B. Cukic, R. Desmarais, S. Dustdar, G. Engels, K. Geihs, K. M. Göschka, A. Gorla, V. Grassi, P. Inverardi, G. Karsai, J. Kramer, A. Lopes, J. Magee, S. Malek, S. Mankovskii, R. Mirandola, J. Mylopoulos, O. Nierstrasz, M. Pezzè, C. Prehofer, W. Schäfer, R. Schlichting, D. B. Smith, J. P. Sousa, L. Tahvildari, K. Wong, and J. Wuttke, *Software Engineering for Self-Adaptive Systems: A Second Research Roadmap*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 1–32.

[2] M. Harman, E. Burke, J. A. Clark, and X. Yao, "Dynamic adaptive search based software engineering," in *Proceedings of the 2012 ACM-IEEE International Symposium on Empirical Software Engineering and Measurement*, Sept 2012, pp. 1–8.

[3] Apache Software Foundation, "Apache tomcat." [Online]. Available: http://tomcat.apache.org/

[4] A. J. Ramirez, D. B. Knoester, B. H. C. Cheng, and P. K. McKinley, "Plato: a genetic algorithm approach to run-time reconfiguration inautonomic computing systems," *Cluster Computing*, vol. 14, no. 3, pp. 229–244, 2011.

[5] S. S. Andrade and R. J. d. A. Macdo, "A search-based approach for architectural design of feedback control concerns in self-adaptive systems," in *2013 IEEE 7th International Conference on Self-Adaptive and Self-Organizing Systems*, Sept 2013, pp. 61–70.

[6] D. El Kateb, F. Fouquet, G. Nain, J. A. Meira, M. Ackerman, and Y. Le Traon, "Generic cloud platform multi-objective optimization leveraging models@run.time," in *Proceedings of the 29th Annual ACM Symposium on Applied Computing*, ser. SAC '14. New York, NY, USA: ACM, 2014, pp. 343–350.

[7] Z. I. M. Yusoh and M. Tang, "Composite saas placement and resource optimization in cloud computing using evolutionary algorithms," in *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*, June 2012, pp. 590–597.

[8] A. S. Sayyad, J. Ingram, T. Menzies, and H. Ammar, "Scalable product line configuration: A straw to break the camel's back," in *Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on*, Nov 2013, pp. 465–474.

[9] N. Esfahani, A. Elkhodary, and S. Malek, "A learning-based framework for engineering feature-oriented self-adaptive software systems," *IEEE Transactions on Software Engineering*, vol. 39, no. 11, pp. 1467–1493, Nov 2013.

[10] G. G. Pascual, R. E. Lopez-Herrejon, M. Pinto, L. Fuentes, and A. Egyed, "Applying multiobjective evolutionary algorithms to dynamic software product lines for reconfiguring mobile applications," *Journal of Systems and Software*, vol. 103, pp. 392 – 411, 2015.

[11] K. Kang, S. Cohen, J. Hess, W. Novak, and A. Peterson, "Feature-oriented domain analysis (foda) feasibility study," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU/SEI-90-TR-021, 1990. [Online]. Available: http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=11231

[12] E. M. Fredericks, "Automatically hardening a self-adaptive system against uncertainty," in *Proceedings of the 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS '16. New York, NY, USA: ACM, 2016, pp. 16–27.

[13] V. Cardellini, E. Casalicchio, V. Grassi, S. Iannucci, F. L. Presti, and R. Mirandola, "Moses: A framework for qos driven runtime adaptation of service-oriented systems," *IEEE Transactions on Software Engineering*, vol. 38, no. 5, pp. 1138–1159, Sept 2012.

[14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *Trans. Evol. Comp*, vol. 6, no. 2, pp. 182–197, April 2002.

[15] Q. Zhang and H. Li, "MOEA/D: A multiobjective evolutionary algorithm based on decomposition," *IEEE Trans. Evolutionary Computation*, vol. 11, no. 6, pp. 712–731, 2007.

[16] K. Li, Q. Zhang, S. Kwong, M. Li, and R. Wang, "Stable matching-based selection in evolutionary multiobjective optimization," *IEEE Trans. Evolutionary Computation*, vol. 18, no. 6, pp. 909–923, 2014.

[17] Rice University, "Rice university bidding systems." [Online]. Available: http://rubis.ow2.org/

[18] M. Arlitt and T. Jin, "A workload characterization study of the 1998 world cup web site," *Netwrk. Mag. of Global Internetwkg.*, vol. 14, no. 3, pp. 30–37, May 2000.

[19] K. Deb and D. Kalyanmoy, *Multi-Objective Optimization Using Evolutionary Algorithms*. New York, NY, USA: John Wiley & Sons, Inc., 2001.

[20] E. Zitzler and S. Künzli, "Indicator-based selection in multiobjective search," in *Proc. of PPSN VIII: the 8th International conference on Parallel Problem Solving from Nature*, 2004, pp. 832–842.

[21] E. Zitzler and L. Thiele, "Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach," *Evolutionary Computation, IEEE Transactions on*, vol. 3, no. 4, pp. 257–271, Nov 1999.

[22] T. Chen and R. Bahsoon, "Self-adaptive and sensitivity-aware qos modeling for the cloud," in *Proceedings of the 8th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2013, pp. 43–52.

[23] T. Chen, R. Bahsoon, and X. Yao, "Online qos modeling in the cloud: A hybrid and adaptive multi-learners approach," in *Proceedings of the IEEE/ACM 7th International Conference on Utility and Cloud Computing*, 2014, pp. 327–336.

[24] D. Benavides, S. Segura, and A. Ruiz-Corts, "Automated analysis of feature models 20 years later: A literature review," *Information Systems*, vol. 35, no. 6, pp. 615 – 636, 2010.

[25] K. Miettinen, *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, 1999, vol. 12.

[26] I. Das and J. E. Dennis, "Normal-boundary intersection: A new method for generating the pareto surface in nonlinear multicriteria optimization problems," *SIAM J. on Optimization*, vol. 8, no. 3, pp. 631–657, 1998.

[27] N. Roy, A. Dubey, A. Gokhale, and L. Dowdy, "A capacity planning process for performance assurance of component-based distributed systems," in *Proceedings of the 2Nd ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '11. New York, NY, USA: ACM, 2011, pp. 259–270.

[28] F. Fittkau, S. Frey, and W. Hasselbring, "Cdosim: Simulating cloud deployment options for software migration support," in *2012 IEEE 6th International Workshop on the Maintenance and Evolution of Service-Oriented and Cloud-Based Systems (MESOCA)*, Sept 2012, pp. 37–46.

[29] University of Cambridge Computer Laboratory, "Xen: a virtual machine monitor." [Online]. Available: http://www.xenproject.org/

[30] Oracle Corporation, "Mysql." [Online]. Available: https://www.mysql.com/

[31] Apache Software Foundation, "Ehcache." [Online]. Available: http://www.ehcache.org/

[32] FEMOSAA, "source code of femosaa, the applied adaptable software and setups of experiments." [Online]. Available: https://github.com/taochen/ssase

[33] MOEA/D-STM, "source code of moea/d-stm with dependency aware operators and knee selection." [Online]. Available: https://github.com/JerryI00/Software-Adaptive-System

[34] A. Bourdon, A. Noureddine, R. Rouvoy, and L. Seinturier, "PowerAPI: A Software Library to Monitor the Energy Consumed at the Process-Level," *ERCIM News*, vol. 92, pp. 43–44, Jan. 2013.

[35] S. Wang, S. Ali, T. Yue, Y. Li, and M. Liaaen, "A practical guide to select quality indicators for assessing pareto-based search algorithms in search-based software engineering," in *Proceedings of the 38th International Conference on Software Engineering*, ser. ICSE '16. New York, NY, USA: ACM, 2016, pp. 631–642.

[36] V. B. Kampenes, T. Dyb, J. E. Hannay, and D. I. Sjberg, "A systematic review of effect size in software engineering experiments," *Information and Software Technology*, vol. 49, no. 1112, pp. 1073 – 1086, 2007.

[37] N. Esfahani, E. Kouroshfar, and S. Malek, "Taming uncertainty in self-adaptive software," in *Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering*, ser. ESEC/FSE '11. New York, NY, USA: ACM, 2011, pp. 234–244.

[38] S. Frey, F. Fittkau, and W. Hasselbring, "Search-based genetic optimization for deployment and reconfiguration of software in the cloud," in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. Piscataway, NJ, USA: IEEE Press, 2013, pp. 512–521.